



Intel[®] Technology Journal

Second-Generation Intel[®] Centrino[™] Mobile Technology

Interface Material Selection and a Thermal Management Technique in Second-Generation Platforms Built on Intel[®] Centrino[™] Mobile Technology

Interface Material Selection and a Thermal Management Technique in Second-Generation Platforms Built on Intel[®] Centrino[™] Mobile Technology

Eric C. Samson, Mobility Group, Intel Corporation
Sridhar V. Machiroutu, Mobility Group, Intel Corporation
Je-Young Chang, Technology and Manufacturing Group, Intel Corporation
Ishmael Santos, Mobility Group, Intel Corporation
Jim Hermerding, Mobility Group, Intel Corporation
Ashay Dani, Technology and Manufacturing Group, Intel Corporation
Ravi Prasher, Technology and Manufacturing Group, Intel Corporation
David W. Song, Technology and Manufacturing Group, Intel Corporation

Index words: cooling, TIM, thermal management, Intel[®] 915, Pentium[®] M processor

ABSTRACT

The mobile thin and light platform has a limited cooling capability, in part due to a form factor that limits the volume available for the thermal solution. The high performance of the Pentium[®] M processor on 90nm process technology and the Intel[®] 915 Chipset in the second-generation platform built on Intel[®] Centrino[™] mobile technology demands a high electrical power and generates substantial heat, presenting a challenge to the thin and light notebook system designer. In this paper, we address two methods of dealing with the thermal challenge.

First, we discuss the path-finding effort to improve the thermal interface materials (TIMs) that allow a good thermal contact between processor and thermal solution, minimizing the transistor temperature of the bare-die Pentium M processors. Two tester methodologies are described, and the need to test TIMs under mobile usage conditions is emphasized. We also discuss the reliability test methodology for TIMs with a focus on the effect of mobile usage conditions affecting long-term reliability of the TIM.

We then focus on power-based thermal state estimation as a platform thermal management technique. This technique is used to detect and limit the thermal impact of power virus¹ workloads. In the second-generation platforms built on Intel Centrino mobile technology, the Intel 915 Chipset Graphics and Memory Controller Hub (GMCH) is uniquely positioned to understand much of the workload for the platform. The Intel 915 GMCH has implemented filter-based thermal management. Several key usage models for filter-based thermal management are explored in detail: detecting and limiting the impact of power viruses on system memory and detecting and limiting the impact of power viruses on chipset memory controller hubs.

INTRODUCTION

Notebook system designs vary significantly from designer to designer; however, they are all densely packed with components and devices, which leave the system with little room for cooling. The problem is compounded by the limited room inside a thin and light notebook product, which typically has a one-inch total thickness when folded and a 17 mm inner vertical space in the lower half of the notebook computer. Figure 1 shows a schematic electrical layout of the major components in the second-generation

[®] Pentium is a registered trademark of Intel Corporation or its subsidiaries in the United States and other countries.

[®] Intel is a registered trademark of Intel Corporation or its subsidiaries in the United States and other countries.

[™] Centrino is a trademark of Intel Corporation or its subsidiaries in the United States and other countries.

¹ A power virus is an unusually intensive workload that maximizes power consumption. Most useful applications draw only a fraction of the power a power virus consumes.

platforms built on Intel Centrino mobile technology. Figure 2 shows a layout of platform-based notebook system that is roughly representative of performance thin and light notebook designs in the industry. In general, the low-profile thin and light form factor limits the flexibility of thermal solution choices for the system components that must be cooled in order to get any appreciable performance.




Figure 1: Electrical schematic of the second-generation platforms built on Intel Centrino mobile technology

Figure 3 shows the use of the remote heat exchange, the predominant method of cooling of high-power components that require dedicated active cooling. In remote heat exchange, the thermal energy is transported to a location, typically via a heat pipe, where a larger fan and heat exchanger can be used. Also shown in Figure 3 are the silicon portion of the hot component (bare die assumed), the attached hardware for coupling the thermal solution to the hot component, and the key temperature monitor points typically used to characterize the performance of the solution. In the first section, we discuss the ability to transfer the thermal energy from the processor to the thermal solution by using thermal interface materials (TIMs).




Figure 2: Layout of typical thin and light notebook (base only with top surface removed)




Figure 3: Remote heat exchanger using a heat pipe embedded in a block

The thermal solution shown in Figure 3 is used to cool the processor, and the fan allows for the cooling of the platform by pulling air flow over various components in the platform. Under thermally high-stress applications such as a power virus or due to improper design of the thermal solution, the platform components could generate a severe thermal environment. To protect the notebook computer and the component functionality, Intel has built multiple thermal management elements into the second-generation platform built on Intel Centrino mobile technology.

THERMAL INTERFACE MATERIALS

For the bare-die Pentium M processor, a successful thermal solution design would allow a minimal temperature drop from the silicon transistor temperature to the ambient temperature. A remote heat exchanger (RHE) is used as the thermal solution for cooling the Pentium M processors in the second-generation platforms built on Intel Centrino mobile technology for thin and light systems. Figure 3 shows a schematic of an RHE. The efficiency of heat draw by the heat pipe (in the attach

block) from the processor depends on the quality of thermal contact between the attach block and the processor. The lower the thermal contact resistance, the lower the temperature drop from the silicon transistor to the ambient.

Even in a direct contact, the processor and the attach block do not transfer heat efficiently, because the quality of contact between two non-conforming solid surfaces is poor, as shown in Figure 4 (a). To enhance the thermal contact between the processor and the attach block, TIMs are inserted into the interface, as shown in Figure 4 (b). Under mechanical pressure, the soft TIMs conform to the microscopic surface contours of the adjacent solid surfaces and increase the (microscopic) area of contact between the thermal solution surface (block) and the silicon die surface (processor) and therefore reduce the temperature drop across this contact.




Figure 4: Magnified schematic of contact between Si surface and Solution surface

The quality of contact between the processor and the attach block, or TIM performance, depends on the quality of the thermal conduction through the TIMs and the quality of contact between the TIMs and the two surfaces. Mathematically this can be represented as follows [1]:

$$\theta_{TIM} = \frac{BLT}{k_{TIM}} + R_c \quad \text{Equation 1}$$

where θ_{TIM} is the effective performance of the TIM, k_{TIM} is the bulk thermal conductivity of the TIM, BLT (Bond Line Thickness) is the thickness of the TIM under usage, and R_c is the contact resistance between the TIM and the mating surfaces. The contact resistance is mainly due to the irregularities or roughness of the surfaces of processor and attach block, so the resistance is negligible, if the surfaces are perfectly smooth.

Based on Equation 1, we can consider three approaches (or strategies) for reducing θ_{TIM} . One approach is to increase the conductivity of TIM, k_{TIM} . This is generally done by using a high thermal conductivity material (like metal or liquid metal) or a low thermal conductivity base material loaded with highly conductive particles. Another approach to reduce θ_{TIM} is to reduce the BLT of the TIM. This is done by reducing the bulk modulus of elasticity of the TIM. The final approach is to reduce R_c by filling the crevices of the processor and attach block surfaces.

Wetting materials allow for low R_c . TIM developers are trying to achieve good performance by optimizing or improving one or more of the three parameters. Table 1 lists various types of TIM, their properties, advantages and issues.

Table 1: Properties and applicability of TIMs

Material Type	Typical composition	Advantages	Disadvantages	+BLT (mil)	Thermal conductivity (W/m-K)
Grease	AlN1, Ag, ZnO, Silicon oil	High bulk thermal conductivity, conforms to surface irregularities, no cure needed, re-usable	Pump out and phase separation (post rel test), migration	2.0	3 to 5
Gel	Al, Ag, Silicone oil, Olefin	Good bulk thermal conductivity, conforms to surface irregularities before cure, no pump out or migration, re-usable	Cure needed, lower thermal conductivity than grease	1-1.5	3 to 4
*PCM	Polyolefins, epoxies, polyesters, acrylics, BN, Alumina, Al, Carbon nanotubes	Conforms to surface irregularities, no cure needed, no delamination, easy handling, re-usable	lower thermal conductivity than grease, no uniform BLT	1.5-2	0.5 to 5
**PCMA	Pure In, In/AG, Sn/Ag/Cu, In/Sn/Bi	High thermal conductivity, easy handling, re-usable	Complete melt possible, voiding	2-5	30 to 50
Solder	Pure In, In/AG, Sn/Ag/Cu, In/Sn/Bi	High thermal conductivity, easy handling, no pump out	Reflow needed, stress crack/delamination (post rel), voiding possible, Not re-usable	2-5	30 to 50

*PCM: Phase Change Material

**PCMA: Phase Change Metallic Alloy

+BLT: Bond Line Thickness of TIM

In addition to improving the thermal performance of contact between two surfaces, TIM is also required to be reworkable; i.e., allow the separation of the processor and thermal solution. Processors are expensive (> \$100) in comparison to the thermal solution (~ \$10) or the TIM (~\$1). If the thermal solution fails or the TIM is incorrectly applied, it should be possible to salvage the processor.

Table 2 lists the performance of best-in-class measurements of TIM, by material type. It also indicates

the potential performance of that material class in the near future.

Table 2: Representative θ_{TIM} performance range of best-in-class TIM (Unit in $^{\circ}\text{C}\cdot\text{cm}^2/\text{W}$)

Material Type	1995	1998	2001	2004	Beyond
Elastomer	1.0	0.7	0.6	0.5	0.4 ?
Grease	0.3	0.2	0.2	0.1	0.05 ?
*PCM	-	0.25	0.2	0.1	0.05 ?
+PCMA	-	-	0.1	0.05	0.03 ?
Gel	-	0.3	0.2	0.08	0.05 ?

*PCM: Phase Change Material

+PCMA: Phase Change Metallic Alloy

Gels (since 2000) and phase change metal alloys (since 2001) are newer to the market. A review of Table 2 also indicates that the potential for improved TIM performance exists based on past trends. Realizing this potential improvement requires the updated formulations of TIMs, the optimization of current each ingredient material of TIMs, the selection of base polymers, and the improvement of filler properties. Since the potential performance improvement diminishes with material maturity (the law of diminishing returns), the tailored optimization of TIMs to a specific application becomes important. The first step in the development of application-specific materials is the ability to quantify the TIM performance under the required application (in this case, notebook systems).

TIM Characterization Methodology


Two methods to quantify the performance of TIMs are discussed next. The first method, often used by material developers, is to quantify the material characteristics of the TIM. This method involves the use of a material tester, and is described first.

Material Tester


In simple terms, a material tester (ASTM D-5470 based [2]) consists of a TIM filled between two blocks with coplanar surfaces. One block is heated and the other block is cooled. The temperature at the interface of each block is measured. The heat flux through the TIM is measured and the TIM resistance (a measure of TIM performance) is calculated as follows:

$$\theta_{TIM} = \frac{T_{hot,int} - T_{cold,int}}{P} \quad \text{Equation 2}$$

Where P is the power (heat flux) conducted through the TIM, $T_{hot,int}$ is the temperature of the TIM interface on the hot side and $T_{cold,int}$ is the temperature of the TIM interface at the cold side. Figure 5(a) shows a schematic representation using the temperature gradient approach, where θ_{TIM} can be calculated from Equation 2.



(a) Schematic showing temperature gradient method to measure θ_{TIM}



(b) Schematic of Material Tester

Figure 5: ASTM-based Material Tester

Figure 5(b) shows a schematic of an ASTM-based Material Tester. It consists of two copper rods, one of which is heated by electric resistive heaters and the other is cooled by a water-cooled thermoelectric chiller; and a facility to apply a mechanical pressure and adjust planarity at the interface between the copper rods. The TIM to be tested is placed between the copper rods. Each copper rod has three thermocouples embedded along its axis to measure the temperature gradient along the copper rod. The BLT is measured by reflecting two laser beams off the hot and cold copper bars, respectively, and measuring the distance of the reflected laser beams both with and without a TIM sample in the material tester. The

increase in the measured distance with the TIM is taken as its BLT.

Table 3 shows the TIM performance and BLT measured in an ASTM-based material tester. The test data range from 0.02 °C-cm²/W to 0.2 °C-cm²/W. The above method is used by many TIM manufacturers to estimate the performance of TIM and to develop new formulations of TIMs. Some advantages of the material tester include the controlled co-planarity between surface of interest, the ability to measure the TIM thickness, the ability to apply an accurate pressure on the TIM, and a controlled uniform temperature at the TIM-surface interface.

Table 3: TIM performance data from Material Tester measured at 90 psi loading on TIM

Material	θ_{TIM} (°C-cm ² /W)	[†] BLT (mil)
Grease 1	0.02	0.3
Grease 2	0.06	0.7
Grease 3	0.16	0.3
Gel 1	0.05	0.3
*PCM 1	0.08	1.1
*PCM 2	0.11	1.9
*PCM 3	0.11	0.3
*PCM 4	0.18	4
Foil 1	0.16	2
Foil 2	0.11	0.4
Solder 1	0.11	5
Solder 2	0.18	5

*PCM: Phase Change Material

[†]BLT: Bond Line Thickness of TIM

A caveat with the material tester is that it does not reflect a realistic use-condition environment for the TIM. In a notebook system, such ideal conditions and controls don't exist, so a Mobile TIM Tester was developed to measure the TIM performance under more realistic application conditions. This is discussed below.

Mobile TIM Tester

A Mobile TIM Tester was developed to ensure a realistic characterization of TIM performance in a notebook environment. As shown in Figure 6, it consists of a Thermal Test Vehicle (that simulates a real processor), a wide, thin, flat copper plank with symmetric fan-heat sinks (to simulate a mobile thermal solution), and a real mobile mechanical attach to load the copper plate to the test vehicle. A Thermal Test Vehicle is made from the same technology as a real processor and is a thermal "replica" of the actual processor. It consists of heating elements within the silicon to heat the die surface (in a manner similar to that expected in a real processor). In notebook systems, a heat pipe is used to move energy from the processor to the fan-heat exchanger. However, no standard heat pipe is available on the market with calibrated performance. Hence, a flat, wide copper plank

with the symmetric placement of over-sized fan-heat sinks is used to replicate a heat pipe-based thermal solution. The properties of the copper plank and the fan-heat sink unit are controlled.

A Mobile Mechanical Attach consists of a dimple plate that applies a point center load on the top of the copper plank to ensure minimal tilt and uniform pressure on the TIM and die surface.




Figure 6: Schematic of Mobile TIM Tester

The temperature of the Thermal Test Vehicle, T_j , is measured based on a thermal sensor on the Thermal Test Vehicle. The temperature on the copper plank, T_p , is read by a thermocouple attached to the center of the copper plank on the top. The power input to the processor, P , is read using a combination of a voltage and current meter, or a power meter.

The performance of the TIM is captured as a resistance from the Test Vehicle to the copper plank, thus:

$$\theta_{j-p} = \frac{(T_j - T_p)}{P} \quad \text{Equation 3}$$

The performance of the TIM in a Mobile TIM Tester measured as θ_{j-p} includes the resistances from the processor and copper plank in addition to the TIM itself.

To obtain the true TIM performance, θ_{TIM} , from a Mobile TIM Tester data, θ_{j-p} , a decoder (translator) is needed that removes the contribution from the processor and copper plank in the θ_{j-p} measurement. For the development of translator (θ_{j-p} to θ_{TIM}), a thermal model (simulation) of the Mobile TIM Tester was developed with two fan-heat sink units, copper plate, TIM, and the processor with PCB. Variation in the TIM performance was correlated to the measured values of θ_{j-p} in the thermal model. This led to the translator below:

$$\theta_{TIM} = 1.026 \theta_{j-p} - 0.095 \quad \text{Equation 4}$$

The above correlation is useful to calculate the TIM performance from measured θ_{j-p} values from Intel's Mobile TIM Tester. This equation depends on the specific Mobile Test hardware. A TIM user is expected to develop an independent correlation to translate measurement to TIM performance. The measured TIM performance values in Intel's Mobile TIM Tester are presented in Table 4 (the

same materials were tested in an ASTM-based Material Tester and are presented in Table 3).

Table 4: TIM performance measurements in Intel's Mobile TIM Tester (Unit in °C-cm²/W)

Material	θ_{j-p}	θ_{TIM}
Grease 1	0.17	0.08
Grease 2	0.19	0.10
Grease 3	0.36	0.26
Gel 1	0.19	0.10
*PCM 1	0.21	0.12
*PCM 2	0.21	0.12
*PCM 3	0.21	0.12
*PCM 4	0.25	0.16
Foil 1	0.22	0.13
Foil 2	0.23	0.14
Solder 1	0.17	0.08
Solder 2	0.23	0.14

*PCM: Phase Change Material

Comparisons of TIM Measurements

For seeking the best-in-class TIM for Pentium M processors, Intel explored various materials. Intel procured the samples of TIMs from more than 10 TIM suppliers. All test samples were tested at least four times for repeatability and an average of the TIM performance data for each TIM was calculated. This test data are presented in Tables 3 and 4 and plotted for comparisons in Figure 7.




Figure 7: Test data comparing Material Tester and Mobile TIM Tester results

The test data in Figure 7 do not indicate a good match between the Material Tester and Mobile TIM Tester (a Mobile TIM Tester duplicates a notebook application environment). The TIM performance in a Material Tester shows a lower thermal resistance generally in comparison to a Mobile TIM Tester, which suggests that the Material Tester generates an optimistic view of TIM performance,

a view that is not achieved in a real notebook environment.

TIM performance data collected in a Material Tester are unlikely to indicate a performance that is representative of a real application. Intel emphasizes the need to test TIMs in real notebook environments to characterize the TIM performance and to make an accurate assessment of the performance of the thermal solution. The method described in the previous section enables the accurate measurement of TIM performance in a notebook system environment without requiring the assembly or testing of an actual notebook.

TIM Degradation

The performance of TIMs can degrade with usage. Degradation in TIM performance depends on usage temperature, the time of usage, mechanical loading, and material properties. Since all factors are not well understood, TIM degradation is characterized using empirical methods. TIM degradation can vary with application and test conditions and is measured in a Mobile TIM Tester environment. The idea, again, is to replicate a true notebook environment and characterize degradation therein (via accelerated testing).

TIM Degradation Estimation

The degradation of TIM performance is measured as a change in the measured value of θ_{j-p} (since θ_{j-p} directly relates to TIM performance, as shown previously in Equation 3) after each stress cycle (e.g., high-temperature bake or temperature-humidity) compared to pre-stress state.

Since testing for long-term (one to five years) degradation requires a long period of time that is not realistic in a lab environment, the reliability characterization process uses accelerated testing. In accelerated testing, the severity of stress generated by critical parameters affecting TIM performance is increased. TIM performance data are collected as a time series under a severe stress environment, and these data are empirically modeled with the stress applied and the period of application of the stress.

For example, a TIM in a notebook system might experience temperatures in the 50-90 °C range over its life of three or more years. Under accelerated testing, the thermal stress is enhanced and the TIM is tested in the 125 to 150 °C range for shorter durations. The Mobile TIM Tester is put in a bake chamber and heated at a temperature of 125 °C or 150 °C for about 2000 hours. The TIM performance is measured at every interval as needed (e.g., 100 hours, 500 hours). In Figure 8, the measured TIM performance values of a sample TIM at three different bake temperatures are plotted as a function of bake time. An empirical model is developed that best

fits the data plotted in the same figure. An example of empirical correlation is a modified Arrhenius model shown below [3]:

$$\theta_{j-p,t} = \theta_{j-p,t=0} + B e^{-\left(\frac{E_a}{KT}\right)} \sqrt{t} \quad \text{Equation 5}$$

where $\theta_{j-p,t}$ is the degraded TIM performance at time, t ; $\theta_{j-p,t=0}$ is the TIM performance (before application of stress at time = 0), B is an acceleration coefficient, E_a is the activation energy for the TIM in Mobile Attach; K is Boltzmann constant, and T is the temperature (in Kelvin scale).




Figure 8: Plot showing change in a particular TIM performance at different stress levels

The rate of degradation of TIM depends on the TIM and its interaction with the Mobile TIM Tester (real notebook environment). Different materials have different degradation rates under similar stress conditions. Once a model is developed, the predictions of TIM performance are extrapolated to the end-of-life condition, based on the stress condition in real mobile usage conditions. The end-of-life represents the Intel-specified service life of a notebook product during which no field failure is allowed.

Different materials have different degradation rates under similar stress conditions. Figure 9 shows an example of TIM performance degradation for three materials: PCM-a, PCM-b, and Gel-a, which are selected as the best choices for Pentium M processors through extensive path-finding efforts with TIM manufacturers. The rate of degradation of Gel-a and PCM-b are very similar but substantially lower than PCM-a. The figure also indicates the impact of long-term stress on TIM performance degradation. After five years of usage, the TIM performance should depend on the long-term usage temperature in a notebook system. The long-term usage temperature depends on the notebook thermal solution design and the usage of

notebook. At lower temperatures, the rate of degradation must be lower.




Figure 9: Projection of TIM performance degradation

The degradation in TIM performance is attributed to a prolonged usage under thermal stresses. However, how degradation occurs is not clearly understood. The analysis of TIM samples in a Mobile TIM Tester subjected to accelerated testing has been performed. Some samples have been sheared and laser cut to view the TIM-to-processor and TIM-to-copper plank interfaces. Delamination of the TIM from either the processor surface or copper plank was observed to start at the corners of the processor and progress toward its center with prolonged accelerated testing (or prolonged usage). So delamination of TIM is speculated to cause the degradation of TIM performance. Changes in the material properties with a prolonged use, evaporation/loss of volatile components, the interaction between the thermo-mechanical elements of the Mobile Attach, and oxidation of TIM are possible reasons.

System designers would like to measure the degradation rate of different TIMs and select a material that has a lower rate of degradation. It is also expected that they would estimate the long-term temperature of the processor and the TIM in their system to capture the degradation and TIM performance accurately.

FILTER-BASED THERMAL STATE ESTIMATION AND MANAGEMENT

This section of the paper is devoted to a method for managing high power consumption.

Figure 1 shows a high-level view of the main silicon components on a second-generation platform built on Intel Centrino mobile technology. All data transfers to and from the system memory are managed by the GMCH. If an application (running on the CPU) wants to get data from the hard disk or network, these data will make their way

from the appropriate port on the ICH to system memory and then from system memory to the CPU. If an application is operating on a large data set, the system memory will also get accessed to service the CPU cache misses. In integrated graphics mode, graphics and display also have high bandwidth data streams to the system memory.

In summary, if the GMCH/system memory is idle, the platform itself is generally in an idle state. If the GMCH/system memory is not idle, it is consuming more power, and the GMCH and memory components are heating up.

However, the GMCH and memory components react differently to the same activity. If the GMCH is performing a memory read, the data input buffers on the GMCH will be consuming power, but the GMCH data output buffers will not be active. From the memory's viewpoint, the GMCH memory read will activate the memory's output buffers, as well as the memory's core. So while the GMCH and memory power consumption and temperatures will be somewhat correlated, they will not be identical. Therefore independent mechanisms are required to detect GMCH and memory overheating.

The CPU has a thermal sensor, and there may be skin temperature monitoring thermal sensors, but usually the rest of the platform components do not have thermal sensors.

Thermal design guidelines are provided to customers to ensure that component over-temperature conditions do not occur for normal workloads. It may be possible, however, to design atypical (power virus) workloads that may cause overheating in some components.

A lumped model of component power is as follows:

$$\text{Power} = \text{Dynamic Power} + \text{Static Power} + \text{Leakage} \quad \text{Equation 6}$$

where

Dynamic Power = $C \times V^2 \times f \times AF$
 Static Power = $I \times V$
 Leakage Power = Leakage \times stacking factor $\times V$
 C is the total capacitance toggling at rate $AF \times f$
 f is the effective operating frequency
 AF is activity factor
 V is voltage change or applied
 I is effective constant current sources
 Leakage represents all the component transistor leakage sources
 Stacking factor takes into account the data-dependent component of Leakage

The expansion of the lumped model of dynamic power is as follows:

$$C \times V^2 \times f \times AF = \sum_i C_i \times f_i \times AF_i \times V_i^2 = \sum_i w_i \times AF_i$$

Equation 7

where each component of die functionality contributes its specific fraction w_i (weight) to the total dynamic power.

The relationship between component die junction temperature and steady state power is as follows:

$$T_{\text{die_junction}} = T_{\text{ambient}} + \Theta_{ja} \times \text{Power} \quad \text{Equation 8}$$

where T_{ambient} is the component ambient temperature, Θ_{ja} is the thermal resistivity ($^{\circ}\text{C}/\text{Watt}$), and Power is the component steady state power.

Equation 8 applies to static conditions only. Workloads are time variant, and so die temperature is also time variant.

Lumped thermal analysis [4] (i.e., assume that thermal energy is leaving the die from all elements on its surface and the temperature of the die is uniform) leads to the time behavior response of temperature to a step response ΔT in temperature as follows:

$$T(t) = \Delta T(1 - e^{-t/\tau}) + \text{constant} \quad \text{Equation 9}$$

This equation corresponds to that of a first order low pass filter, and can be written recursively as

$$T(n) = (1 - \alpha)T(n-1) + \alpha \Delta T + \text{constant} \quad \text{Equation 10}$$

where n is the time index, normalized to a sampling frequency of 1 and α is the filter time constant.

We can then combine Equations 8, 9, and 11 as

$$T(n) = (1 - \alpha)T(n-1) + \alpha \Theta_{ja} \sum w_i AF_i + \text{constant} \quad \text{Equation 11}$$

This equation is that of a weighted input power filter.

It has been found (empirically) that Equation 12

$$T(n) = \sum \Delta T_i (1 - e^{-t/\tau_i}) + \text{constant} \quad \text{Equation 12}$$

gives a good fit to the time variant behavior of the die junction temperature to a step response in power (the constant term is a function of I, stacking factor, and ambient temperature). For the current generation silicon fabrication process, the leakage and constant current source of GMCH, ICH, and memory power will tend to be relatively constant at the higher power levels. So measurements of component AF_i can be converted into a reasonable indicator of component power, which can in turn be converted into an indication of die junction temperature.

The maximum possible power will correspond to states that toggle on every clock ($AF_i = 1$). But this rate of

toggling is unlikely to happen for the total capacitance of a component. For example, the very highest activity factors are found in clock trees.

So not only is the activity factor proportional to power, high AF_i serves as an indicator of workloads that are of relatively low utility². It makes sense then to combine information about AF_i to slow down the GMCH core and/or memory activity when over-temperature conditions are suspected:

While $T(n) > \text{Threshold_Temp}$ ³, throttle Equation 13

The next section details the implementation of GMCH and memory throttling mechanisms in the Intel 915 GMCH.

SYSTEM MEMORY AND GMCH THROTTLING USAGE MODELS

The Intel 915 GMCH has two independent mechanisms that cause system memory bandwidth throttling: GMCH thermal management, and DRAM thermal management.

GMCH Thermal Management

GMCH thermal management ensures the GMCH chipset is operating within thermal limits. The underlying theory is that GMCH heating is caused by the activity required to access system memory. The implementation provides a mechanism that controls the amount of system memory activity (i.e., Double-Data-Rate-II (DDR2) IO transactions) to a programmable threshold limit as per Equation 13. Memory activity throttling blocks all transactions or a selected set of transactions to the system memory.

System Memory

System DRAM thermal management ensures that the DRAM modules are operating within thermal limits. The DRAM modules are not equipped with thermal sensors, so their temperatures must be tracked via indirect means. The underlying theory is that a DRAM device heats up by different amounts based on the type of activity it is subject to. For example, the amount of heat contributed by a read command is different to that contributed by a write command. The implementation accounts for this variation by using the appropriate values for w_i for each memory transaction type. Throttling can be initiated by a DRAM activity measurement exceeding a programmed threshold.

² There are exceptions such as test patterns for silicon validation, for example.

³ Threshold_Temp is set taking into account product reliability limits, and sufficient guardband to compensate for the estimation error in $T(n)$.

GMCH Thermal Throttling

If the weighted power transaction filter output exceeds a programmable activity threshold (Threshold_Temp_GMCH), then the GMCH starts throttling GMCH activity. As per Equation 13, the throttling lasts for as long as the throttling threshold is exceeded.

Since GMCH thermal throttling is specific to the GMCH, there are three types of transactions (AF_i) that can be assigned different weights (w_i):

1. *Write*: This value is loaded in the filter when a write command is issued on the bus.
2. *Idle*: This value is loaded in the filter when there is no command being issued on the memory bus.
3. *Non-Write*: This value is loaded in the filter when a command other than write (read, precharge, activate, etc.) is issued on the bus.

Figure 10 shows an example of the relationship between the threshold value chosen for throttling and the GMCH power. Figure 11 shows that 3D graphics performance is unaffected for throttling levels above 38 h (programmable activity threshold). Since this behavior is typical of other workloads, a throttling level of >38 h will filter out undesirable workloads while leaving desirable applications unaffected.




Figure 10: GMCH power as a function of throttling threshold




Figure 11: Relative graphics benchmark as a function of throttling threshold

DRAM Rank-Based Throttling

System DRAM devices are organized in ranks. Typically, the memory devices on a side of a memory module form a rank. Each rank heats up independently based on the activity it is subject to by the GMCH. Hence each rank requires an independent power filter.

If any of the rank weighted power transaction filter outputs exceed a programmable activity threshold (Threshold_Temp_DRAM), then the GMCH starts throttling memory. As per Equation 13, the throttling lasts for as long as the throttling threshold is exceeded.

DRAM devices have many pertinent power states:

1. Active State: Reads with Autoprecharge
2. Writes with Autoprecharge
3. Bust Reads
4. Burst Writes
5. Precharge All
6. Precharge
7. Activate
8. Page Open Idle
9. Page Close Idle (all banks precharge)
10. Page Open Power down
11. Page Close Power down (all banks precharge and power down)
12. Rank in Self refresh

The power values for these states are DRAM vendor specific. The Joint Electron Device Engineering Council (JEDEC) has a Serial Presence Device (SPD) specification⁴ that allows the memory module vendors to

⁴ Joint Electron Device Engineering Council DDR2 SPD Revision 1.1

fill in the register fields for the equivalent temperature delta ($= \Theta_{ja} \times w_i$) that can be mapped into most of these power states. These values can be read by the system bios from the memory modules during boot time and placed into the GMCH registers that store the state weights. After this operation is completed, the DRAM filters will look up and use the appropriate stored values as the particular memory states occurs.

For example, the memory module's SPD has the DT4R register field reserved for the temperature rise from ambient due to continual burst read operation. So each time a new memory read burst is started on a particular rank of memory, the DT4R value is used as an input into the filter for that rank.

Figure 12 shows memory power varying as a function of throttling level (expressed as a function of equivalent read bandwidth) for two different memory vendors.




Figure 12: DRAM power

Since one vendor's power consumption per unit bandwidth is less than another's, the filter output for the lower power vendor will be smaller, and hence less likely to cause an overheating indication.

Thermal Power Filter Usage In Conjunction with Other Methods

While this method is able to monitor power variations due to GMCH activity, there are several factors that need to be encompassed in the guardband applied to the thresholds used to decide throttling.

For example, this method is not aware of the actual ambient temperature, so must assume the worst-case specification. Voltage variation must also assume worst-case specifications.

Since an approximation to the thermal diffusion equation is used, the actual thermal transient behavior modeled will be conservatively set so that it is slower for temperature increases, and faster for temperature decreases. This

mismatch also contributes to the guardbanding requirements.

On-die thermal sensors, when available with sufficient die coverage and accuracy, are a better way to deal with these uncertainties. But since an over-temperature condition may last an indefinite time, it is desirable to allow some work to progress rather than hang the platform (due to halting all activity) until the temperature has fallen back to reasonable levels. The filter-based method is a good way to control the amount of throttling happening during over-temperature conditions since it can allow work to proceed at a reduced rate:

```

While (thermal sensor is sensing overheating) {
    While T(n) > threshold {
        throttle.
    }
}

```

Equation 14

While circuitry for direct power measurements can also be used to obtain instantaneous power measurements, these power measurements would still need to be filtered to correlate with the thermally relevant workload.

SUMMARY

Improving the platform component performance/power efficiency and improving the box cooling capability are the primary vectors to maximizing the platform performance given finite heat budgets. By using better material for component packages, component cooling is improved. Improvements in TIM performance are expected, and better characterization of the performance and reliability of these materials will aid this process.

Data were presented to demonstrate the improved accuracy of a Mobile TIM tester or testing in real-use conditions rather than just ASTM D5470-based Material testing.

A simple, accurate, and accelerated reliability test method based on the Arrhenius reliability model was presented.

Notebook system designers and TIM vendors are encouraged to study the degradation of TIM in real systems using the methodology outlined here.

Once a platform's cooling solution and components have been chosen, there is still some added scope to maximize performance by minimizing the guardband between the operating states allowed and the physical thermal limits.

We have shown how filter-based thermal state estimation can be used to help control overheating in platform components, even when these platform components do not contain internal thermal sensors. We have shown how these filters can be performance-transparent in many usage scenarios, and yet still be able to detect and limit power

viruses. We have also shown how filters can be useful even if there is a thermal sensor available. Filter-based thermal estimation enables the second-generation platform built on Intel Centrino mobile technology to set a higher level of performance than it might otherwise allow, while still providing some protection from power viruses.

ACKNOWLEDGMENTS

We thank Lloyd Pollard and Yuchen Huang for supporting much of the work related to JEDEC standardization of memory event weights, and Terry Fletcher for the GMCH power weight event analysis. We thank Pachi Chen for providing the initial data on thermal transient characterization of GMCH heating, and Abdul Momem for providing the platform performance sensitivity to throttling.

REFERENCES

- [1] R. S. Prasher, J. Shipley, S. Prstic, P. Koning and J. L. Wang, "Thermal Resistance of Particle Laden Polymeric Thermal Interface Materials," in *Proceedings of International Mechanical Engineering Congress and Exposition*, Nov. 15-21, 2003, Washington, D.C.
- [2] ASTM, "Standard Test Method for Thermal Transmission Properties of Thin Thermally Conductive Solid Electrical Insulation Materials," ASTM D-5470-01, 1999.
- [3] F. W. Billmeyer, Jr., *Textbook of Polymer Science*, John Wiley Publication, 3rd Edition, pp. 72-73, 1984.
- [4] D. Pitts and L. Sissom, *Schaum's Outlines Heat Transfer*, 2nd Edition, pp. 88-89, 1998, McGraw-Hill, New York.

AUTHORS' BIOGRAPHIES

Eric C. Samson joined Intel in 1983 to design communication chipsets. He currently focuses on power management architecture for chipsets. His e-mail is eric.c.samson at intel.com.

Sridhar V. Machiroutu joined Intel in 2000 with a Ph.D. in Mechanical Engineering. He works on developing cooling technologies for mobile platforms including TIMs and heat pipes. His e-mail is sridhar.v.machiroutu at intel.com.

Je-Young Chang joined Intel in 2000 with a Ph.D degree in Mechanical Engineering. He worked in the Mobility Group developing thermal solutions for mobile processors and platforms for four years, and then in Assembly Technology Development developing advanced thermal packaging technologies. His e-mail is je-young.chang at intel.com.

Ishmael Santos works on Mechanical/Thermal Enabling in the Mobility Group. His interests include computer architecture with an emphasis on platform thermals, power consumption, and performance. He received his B.S. degree in Electrical Engineering and Computer Science from the University of California, Los Angeles. His e-mail is ishmael.f.santos@intel.com.

Jim Hermerding joined Intel in 1993 with a B.S. degree in Computer Engineering. He has worked in the Mobility Group for 11 years in various roles including architecture, engineering, product technical marketing, and thermal and mechanical applications. His current role is a manager within the Thermal Mechanical Enabling team. This involves managing a team to help define, test, and enable thermal control and protection mechanisms in Intel mobile products and to deliver component thermal testing software. He also manages a team of thermal simulation engineers to analyze cooling limits of mobile platforms. His e-mail is jim.g.hermerding@intel.com.

Ashay Dani obtained his Ph.D degree in Chemical Engineering in 1995. He worked for five years in various roles supporting plastics molding materials and process research at KeyTronic Corp (ex-Honeywell), before joining the Intel Assembly Technology Development Group in 2000. Within Intel, he has supported development of various polymer packaging materials and suppliers (thermal interface materials, thermal heat spreaders, sealants, adhesives, die attach, encapsulants, molding compounds, etc.) for several internal and external packaging platforms and products. He has published several papers in referenced journals and conferences within Intel. His e-mail is ashay.a.dani@intel.com.

Ravi Prasher joined Intel in 1999 with a Ph.D degree in Mechanical Engineering. He has worked on many aspects of TIMs and heat pipes for mobile applications and is currently involved in evaluating microchannel cooling technology for electronics cooling. He currently holds eight US patents with approximately 20 patent applications pending with the US patent office. He has published more than 40 articles in journals and conference proceedings. He is also an adjunct professor in the Department of Mech. and Aerospace Engineering at Arizona State University. His e-mail is ravi.s.prasher@intel.com.

David W. Song joined Intel in 2003 with a Ph.D degree in Mechanical Engineering from UCLA. He is a thermal engineer in the Design Process Development Group in Assembly Technology Development. He is currently working on the batch thermal property characterization of mobile test packages for product certification, as well as the current carrying capacity characterization of socket contacts. His graduate study and continuing research interests are in the microscale heat transfer in

semiconductor and nano/microcomposites. His e-mail is david.w.song@intel.com.

Copyright © Intel Corporation 2005. This publication was downloaded from <http://www.intel.com/>.

Legal notices at <http://www.intel.com/sites/corporate/tradmarx.htm>.

For further information visit:

developer.intel.com/technology/itj/index.htm